

**Massachusetts Institute of Technology
Lincoln Laboratory**
244 Wood Street
Lexington, MA 02420-9108

**Massachusetts Institute of Technology
Computer Science and Artificial Intelligence Laboratory (CSAIL)**
The Stata Center, 32 Vassar Street
Cambridge, MA 02139

**MIT CSAIL and Lincoln Laboratory
Task Force Report**

TBD 2016

Prof. Regina Barzilay (CSAIL), Co-chair
Mr. Robert Bond (MIT LL), Co-chair
Prof. Arvind (CSAIL)
Dr. Michael Boulet (MIT LL)
Dr. Robert Cunningham (MIT LL)
Prof. Srinivas Devadas (CSAIL)
Mr. Ken Gregson (MIT LL)
Dr. Jeremy Kepner (MIT LL)
Dr. Sanjeev Mohindra (MIT LL)
Dr. Hamed Okhravi (MIT LL)
Prof. Daniela Rus (CSAIL)
Prof. Julie Shah (CSAIL)
Dr. Howard Shrobe (CSAIL)
Dr. Michael Vai (MIT LL)
Dr. David Whelihan (MIT LL)
Dr. Beijia Zhang (MIT LL)

This material is based upon work supported under Air Force Contract No. FA8721-05-C-0002 and/or FA8702-15-D-0001. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the U.S. Air Force.

© 2016 MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Delivered to the U.S. Government with Unlimited Rights, as defined in DFARS Part 252.227-7013 or 7014 (Feb 2014). Notwithstanding any copyright notice, U.S. Government rights in this work are defined by DFARS 252.227-7013 or DFARS 252.227-7014 as detailed above. Use of this work other than as specifically authorized by the U.S. Government may violate any copyrights that exist in this work.

This page intentionally left blank.

TABLE OF CONTENTS

	Page
List of Illustrations	v
List of Tables	vii
1. BACKGROUND AND TERMS OF REFERENCE	1
2. EXECUTIVE SUMMARY	3
3. DETAILED DISCUSSIONS, FINDINGS, AND RECOMMENDATIONS	7
3.1 Previous and Existing Collaborations	7
3.2 Collaborations to Enhance CSAIL Undergraduate Education	11
3.3 Collaborations to Facilitate Technology Transfer from CSAIL to MIT Lincoln Laboratory	12
3.4 Collaborations to Create Funded Projects to Address Grand Challenges in CSAIL Domains	13
3.5 Collaborations to Create On-Site Research and Development Opportunities between MIT Lincoln Laboratory and CSAIL	32
4. SUMMARY	33
APPENDIX A. CSAIL DIRECTOR'S INITIATIVE	35
APPENDIX B. MIT BEAVER WORKS CAPSTONE PROJECTS	41
APPENDIX C. MIT LINCOLN LABORATORY INVOLVEMENT IN MIT UNDERGRADUATE PROGRAMS	43
LIST OF ACRONYMS	45

This page intentionally left blank.

LIST OF ILLUSTRATIONS

Figure No.		Page
1	MIT Lincoln Laboratory funding to CSAIL.	8
2	MIT Lincoln Laboratory funding to CSAIL, with the largest two programs removed. This plot removes the larger “outliers” and gives an indication of the variance for nominal programs for the past 8 years. Note: the outliers removed were due to the Agile Robotics program with Prof. Seth Teller.	8
3	Combining strengths of both institutions.	14
4	Continuous resupply for humanitarian assistance and disaster relief.	15
5	Autonomous standoff system architecture.	16
6	Legged autonomous ground vehicle (Jaguar).	17
7	Two orthogonal frameworks for autonomous systems.	18
8	Autonomous architecture components.	20
9	MIT SuperCloud software stack used across MIT Lincoln Laboratory and MIT supercomputers.	24
10	New architecture for low-power off-grid near-sensor processing.	25
11	BigDAWG federated database architecture.	25
12	Secure embedded computing evolution: past, present, and future. Academia dominates the future open and universal security technology development.	27
13	Just-In-Time Secure Thread (JIT-ST) processor performs encryption, decryption, and key management deeply within the processor hardware to enable mutually distrusting software to work closely together without divulging secret data. Attack surface is minimized without significant impact on performance or ease of use.	29
14	Secure application-specific systems’ (SASS) vision; the mission system design space will be expanded to include a dimension in security. The objective is to enable the co-design of functionality and security from the get-go.	31

LIST OF ILLUSTRATIONS **(Continued)**

Figure	No.	Page
15	SASS program concept. A tool set and libraries will be developed so that system developers can determine security requirements by analyzing applications and potential threats. The security requirements become the inputs of the system, which performs automatic or semi-automatic SE/HW security co-design. The result is synthesized (HW) and compiled (SW) into application-specific secure system architecture.	31

LIST OF TABLES

Table No.		Page
1	Partial List of Principal Investigators and Potential Research Collaborations	22
2	Ongoing Collaboration between CSAIL and the Cyber Security and Information Sciences Division	28

This page intentionally left blank.

1. BACKGROUND AND TERMS OF REFERENCE

The MIT Computer Science and Artificial Intelligence Laboratory (CSAIL) and MIT Lincoln Laboratory (MIT LL) are, each on their own, critical contributors to MIT's mission of advancing knowledge, educating students, and serving the nation and the world. Increased collaboration between the two organizations in recent years has enabled significant new accomplishments in education, applied research, and system prototyping. For example, the CSAIL-MIT LL Director's Initiative over the past two years has led to advances in cyber security, low power electronics for database processors, and autonomous systems. We would like to build on these initiatives and other collaborations to create an even stronger and more vibrant relationship between our two organizations.

To achieve this goal, a joint task force was formed, effective 28 April 2015. The task force addressed the following actions:

1. Review, categorize, and report on existing collaborations between CSAIL and MIT LL. Assess current collaborations and make recommendations for potential continuation and enhancement.
2. Assess the role increased collaboration could play in
 - a. Enhancing CSAIL undergraduate education through greater exposure of students to practicing engineers and engineering practice.
 - b. Facilitating MIT LL's use of advanced technology developed by CSAIL researchers.
 - c. Creating funded projects that combine multiple CSAIL faculty with the staff, resources, scale, and technological memory of MIT LL to address grand challenges within the CSAIL research and development domains.
 - d. Creating opportunities for CSAIL faculty, research staff, and students to work on research and prototyping projects part time at MIT LL.
3. Identify specific topics, projects, and actions within each of the categories above to recommend to CSAIL and MIT LL leadership.
4. Report on those recommendations to both organizations.

The Task Force membership included the following participants from MIT LL and CSAIL:

Prof. Regina Barzilay (CSAIL), co-chair
Mr. Robert Bond (MIT LL), co-chair
Prof. Arvind (CSAIL)
Prof. Srini Devadas (CSAIL)

Prof. Daniela Rus (CSAIL)
Dr. Howard Shrobe (CSAIL)
Prof. Julie Shah (CSAIL)
Dr. Michael Boulet (MIT LL)
Dr. Robert Cunningham (MIT LL)
Mr. Ken Gregson (MIT LL)
Dr. Jeremy Kepner (MIT LL)
Dr. Sanjeev Mohindra (MIT LL)
Dr. Hamed Okhravi (MIT LL)
Dr. Michael Vai (MIT LL)
Dr. David Whelihan (MIT LL)
Dr. Beijia Zhang (MIT LL)

2. EXECUTIVE SUMMARY

MIT LL and CSAIL have enjoyed several collaborations over the past decade. Collaborative projects have been very diverse, spanning several areas of CSAIL concentration, including robotics, big data analytics, wireless communications, computing architectures and technologies, and cyber security. Nevertheless, the consensus of the committee is that, given the large overlap in interests and the dramatically growing importance to national security of several key CSAIL core areas, collaborations and interactions not only could, but should, increase significantly.

The main task force findings and recommendations are

1. CSAIL-LL projects account for about 1% of CSAIL funding volume, while comprising 11% of Federal flow-through within CSAIL. Currently, CSAIL accounts for 16% of the volume that MIT LL spends on campus. In 2015, the MIT LL funding to CSAIL was fifth among MIT organizations, after the Research Laboratory of Electronics, the Haystack Observatory, the Mechanical Engineering Department, and the AeroAstro Department. Given the relevance and growing importance of CSAIL research and development activities to MIT LL missions, these findings indicate that the funding may be disproportionately low. Moreover, the growing importance of autonomous systems, cyber security, computing architectures, etc. to DoD mission areas indicates that the alignment to DoD interests should be very strong; supporting the conclusion that funding could be fruitfully increased.
2. Current collaborations involve a small portion of CSAIL Principal Investigators (PIs), many of whom have long-term collaborations with LL researchers. Despite significant interest in collaborations with MIT LL PIs, our survey of the broader CSAIL community revealed that most CSAIL PIs are not aware of mechanisms to get engaged in collaborative projects.
3. Most of the current projects involve internal MIT LL funding. Little progress had been made in obtaining joint external funding. Notable exceptions include (1) joint programs between MIT LL's Human Language Technology Group and CSAIL's Spoken Language Systems Group, which have been ongoing for a decade, and (2) the recent NSA-funded collaborations between the MIT Lincoln Laboratory Supercomputing Center and Professors Stonebraker and Edelman.
4. Both of the successes noted above in (3) employed a working model where either MIT LL staff worked closely on campus with CSAIL researchers or where students are supported by MIT LL and work closely with both CSAIL and MIT LL PIs. This close-proximity model seems to engender the best collaborations.
5. The CSAIL Director's Initiative (Appendix A) has produced some notable research and development advances in the areas of cyber security, low-power electronics for databases, and autonomous

systems. However, this initiative has limited reach into the CSAIL PI community and has not yet resulted in sustained external funding.

6. Our recommendations for future collaboration include

- a. Focus on increasing awareness of CSAIL PIs with regards to project opportunities, announcing the deadlines, and clearly communicating selection criteria.
- b. Continue the CSAIL-MIT LL Director's Initiative for another three years, to help encourage research collaborations in areas of growing national security needs such as cyber and autonomous systems. However, shift the award process to solicit projects for the Director's Initiative via an open submission process. Areas of interest and selection criteria will be decided by a committee nominated by CSAIL and MIT LL leadership. Potential for transition to eternally sponsored program will be one of the key selection criteria.
- c. Encourage MIT LL staff to work with CSAIL researchers, students, and post-docs on campus. Task force members have noted that direct MIT LL presence on the order of one or more days per week has led to strong and growing research collaborations in the past.
- d. Facilitate technology transfer by involving MIT LL researchers in CSAIL projects of strategic interest to MIT LL. Long-term visits of MIT LL visitors are identified as a particularly promising mechanism in this area.
- e. Pursue opportunities for building research connections between MIT LL and CSAIL researchers and identify areas of research interest. Possible venues include seminars, targeted meetings, and long-term visits of MIT LL researchers on campus.
- f. Jointly pursue large-scale externally funded sources. As a step in this direction, the committee proposed four subcommittees to formulate grand challenges in the areas of strategic interest to national security.
- g. Expand CSAIL involvement in the Beaver Works Center.
- h. MIT LL and CSAIL leadership should periodically review and evaluate collaborations and opportunities for increased interaction.

Below we elaborate on these items.

The committee recognizes and appreciates the value of projects funded by the MIT LL Technology Office (TO). Over the years, there have been numerous such projects that have been beneficial to both organizations. The projects have served as valuable conduits for transitioning technology and expertise from CSAIL to MIT LL, while at the same time providing CSAIL with insight into real-world national security problems. However, the number of CSAIL PIs participating in TO projects has remained small.

Also, we note that often the same PI is involved in several projects over the years, and while we recognize that long-standing relationships are certainly valuable, this may also point to an overly entrenched or narrow commerce between the two organizations.

The survey of CSAIL PIs revealed that while many of them are interested in getting actively engaged in the LL collaborations, they are not aware of the avenues for starting such engagements. The perceived “closed club” patterns of current collaborations prevents active engagements on the part of the CSAIL PIs and stifles potential collaborations.

The committee therefore recommends that measures be taken to encourage wider awareness within CSAIL of LL TO project opportunities. Moreover, CSAIL PIs more broadly need to be made aware of the formal mechanisms and timelines for establishing collaboration projects. Indeed, increased mutual awareness of research interests and areas of expertise would be very beneficial. Greater interaction through a CSAIL-MIT LL seminar series and other, informal venues would help. The seminars could be held both on campus and at MIT LL, which would encourage participation from both organizations and would also provide opportunities to get to know each other’s facilities.

The funding level of TO projects is relatively small compared to what is possible with externally funded projects. Hence, the committee recommends the formation of domain-based teams comprising CSAIL and MIT LL PIs to actively pursue larger scale externally funded projects with agencies such as DARPA and IARPA. Towards this end, the committee created four subcommittees to formulate research agendas that address grand challenges, where the combined capabilities of CSAIL and MIT LL could result in significant breakthroughs. We expect these four subcommittees to remain in place after the conclusion of the study, with the goal of pursuing program opportunities in each of the four areas. The four areas examined were

1. Autonomy-in-motion
2. Autonomy-at-rest
3. Advanced computing architectures
4. Cyber security

Challenges, research and development opportunities, and next steps were identified for each of these domains.

The CSAIL Director’s Initiative (see Appendix A), which has been ongoing for the past two years, has generated some strong collaborations between CSAIL and MIT LL, with valuable work emerging in the areas of low-power electronics for databases, functional encryption techniques, just-in-time secure computing, and the application of coresets to autonomous systems. One of the lessons learned from this initiative is the value of constant and close communication between MIT LL and CSAIL PIs, needed to foster technology transition and to keep the goals of each synergistically aligned.

Going forward, the committee recommends that the CSAIL Director's Initiative will be open to a Laboratory-wide proposal submission. The selection process will be conducted by the committee nominated by the MIT LL and CSAIL leadership. The proposed open-selection process will encourage new collaboration between the laboratories, increase awareness about national security problems among CSAIL researchers, and facilitate technology transfer.

MIT LL's access to real-world problems and its emphasis on solution prototyping can help to enhance MIT undergraduate education. CSAIL and MIT LL should seek to develop more capstone projects at the Beaver Works Center, which is dedicated to enhancing research and innovation through project-based learning. We also encourage increasing the number of CSAIL interns at MIT LL through the established IV-A, UPOP, and UROP programs. One novel suggestion is to establish a student mentoring program between MIT LL and CSAIL.

The committee recognizes the need and deep value in establishing a continuous and high volume conduit for technology transfer from CSAIL to MIT LL. CSAIL is a source of tremendous technology research and development. It has world renowned academic researchers working in areas of computer science and artificial intelligence, spanning fundamental research, innovative systems, and advanced technologies. It has spawned dozens of high-tech companies and is a major source of economic expansion in New England and beyond. Moreover, the outstanding work going on at CSAIL has never been more relevant to the many aspects of national security that MIT LL is chartered to address.

In general, greater awareness by MIT LL staff of the work being done on campus would help to prime this conduit, as would greater awareness on the part of CSAIL faculty of the national security problems being addressed at MIT LL. We feel, therefore, the single most effective measure to creating and sustaining a technology pipeline will be to embed MIT LL staff in CSAIL research projects that align with MIT LL interests. This will develop long-term, trustful collaborations with CSAIL, will help CSAIL better achieve its research goals, will allow MIT LL staff to become very familiar with CSAIL technologies and capabilities, will help attract top CSAIL students to MIT LL, and will allow for more effective situational awareness and cultural alignment between the two organizations. Funding for these staff would come from a combination of the TO (initial seeding) and established programs.

We also recommend that leadership in both organizations actively and continuously encourage collaborations and intellectual commerce. Moreover, leadership should meet on a regular basis to review progress and to continue to foster ideas for enhanced collaborations.

3. DETAILED DISCUSSIONS, FINDINGS, AND RECOMMENDATIONS

3.1 PREVIOUS AND EXISTING COLLABORATIONS

MIT LL and CSAIL have enjoyed several collaborations over the past decade. Collaborative projects have been very diverse, spanning several areas of CSAIL concentration, including robotics, big data analytics, wireless communications, computing architectures and technologies, and cyber security. Nevertheless, the committee consensus is that given the large overlap in interests and the emerging importance to national security of several key CSAIL core areas collaborations not only could, but should, increase significantly.

In most cases, collaborations have been funded through MIT LL, most notably through by the MIT LL Technology Office (TO). TO funding has come through a diverse set of venues: ASD(R&E) Line projects, Advanced Concept Committee (ACC) projects, New Technology Initiative (NTI) projects, University Initiatives, Research Assistantships, and the CSAIL Director's Initiative (see Appendix A for a discussion of the latter).

There have also been a few collaborations funded by external sponsors, especially in the area of human language technology. In particular, MIT LL's Human Language Technology (HLT) group has a long-term collaboration that dates back over fifteen years, predating the CSAIL name. Most of this group's focus at CSAIL is within their Spoken Language Systems (SLS) group (<https://groups.csail.mit.edu/sls/>). Jim Glass is a leader of the SLS group and a senior research scientist; the MIT LL HLT group has worked with him for many years. Currently, Jen Drexler is one of the HLT group's LSP (Lincoln Scholar's Program) students doing graduate work in the CSAIL HLT group. The HLT group also supports Stephen Shum, a grad student in the SLS group. Up until recently, the group also supported Najim Dehak, who was a research scientist in the SLS group, but he has now joined Johns Hopkins University. Recently, MIT LL has worked with CSAIL on joint computer database and cloud computing projects funded by the National Security Agency (NSA). These are collaborations between the MIT LL Computing Center, led by Albert Reuther and Jeremy Kepner, and noted computer science experts at CSAIL, including Professors Stonebraker and Edelman. In these cases, the funding has been about \$2M and has been sent directly to CSAIL, with a large portion flowing through CSAIL to fund MIT LL activities.

Funding from MIT LL to CSAIL from 2010 through to 2015 is shown in Figure 1. There were 21 CSAIL principal investigators involved in these collaborative projects, and 54 projects were funded in all (if a project spans N years, the overall project has been counted as N separate projects since general projects must be recompeted each year). In 2016, 11 CSAIL PIs have ongoing projects with MIT LL. The number of CSAIL PIs working on MIT LL projects compared to the total number of CSAIL PIs (which is over 100) has been quite small. Once again, this indicates that there is a large set of resources at CSAIL that remain untapped by MIT LL.




Figure 1. MIT Lincoln Laboratory funding to CSAIL.




Figure 2. MIT Lincoln Laboratory funding to CSAIL, with the largest two programs removed. This plot removes the larger “outliers” and gives an indication of the variance for nominal programs for the past 8 years. Note: the outliers removed were due to the Agile Robotics program with Prof. Seth Teller.

The median project size is \$100K. The average project size is \$205K, significantly larger than the median. This skew is attributable to a small number of very large projects. Figure 2 shows the funding profile from 2010 through 2015 with the two largest projects (“outliers”) removed. With this adjustment, the average approaches the median and becomes \$119K.

Overall, CSAIL-MIT LL projects account for 1% of CSAIL volume, which comprises 11% of Federal Flow Through within CSAIL. Compared to the overall volume of MIT LL funded projects, CSAIL accounts for 16% of the volume spent on campus. In 2015, CSAIL volume is in fifth place after RLE, Haystack Observatory, Mechanical Engineering, and AeroAstro.

The committee also informally interviewed some of the CSAIL faculty involved in these collaborations. A few of the common threads based on the interviews are as follows:

- Both CSAIL and MIT LL researchers are very enthusiastic about the collaborative efforts and have strong technical ability on both sides. There continues to be ongoing discussions about further collaborations, based on the funded efforts, even after the initial projects have been completed.
- Some CSAIL faculty and MIT LL researchers have attempted to write joint larger scale proposals, and there have been some substantial successes so far (for example, as mentioned above with the MIT LL HLT-CSAIL SLS and the LLSC-CSAIL collaborations), but there is the feeling that much more could be accomplished. Some common elements of these successes include:
 - Many years of prior collaboration between the faculty member and the MIT LL staff member.
 - Strong desire by the sponsor to exploit the distinct strengths of CSAIL faculty (world-class ideas, brilliant researchers, open research environment, etc.) and MIT LL staff (world-class systems engineering, research infrastructure, deep understanding of sponsor problems, professional researchers, etc.).
 - MIT LL staff committed to spending substantial time on the MIT campus (two or more days per week). Ultimately, it would be ideal if every MIT LL staff member aimed to spend some dedicated time each month on the MIT campus to enhance their curiosity, increase collaboration, and advance their careers.
 - MIT LL staff who are passionate about advancing the CSAIL faculty members research vision.
- Through these jointly funded efforts, a number of positive outcomes have emerged:

- The ability of CSAIL faculty to leverage MIT LL staff to take on larger, more ambitious research projects while keeping their post-docs and graduate students focused on basic research.
- The ability of CSAIL faculty to leverage MIT LL staff to effectively utilize more MIT undergraduates and CSAIL students in their research.
- The transition of CSAIL faculty innovations to MIT LL staff for application to MIT LL projects.
- The transition of MIT undergraduates and graduate students to MIT LL as hires.
- CSAIL faculty referring potential sponsors to specific MIT LL staff to transition the CSAIL faculty member's work to the sponsors.

The CSAIL Director's Initiative was created especially to encourage longer-term collaborations between MIT LL and CSAIL. Several fruitful projects have been funded, and strong relationships between CSAIL and MIT LL have been developed as a result of this initiative. At the same time, one of the key lessons learned from this initiative is the value of constant and close communication between MIT LL and CSAIL PIs, needed to foster technology transition and to keep the goals of each synergistically aligned.

The committee conducted an online survey of CSAIL PIs to assess their interest in collaborations with MIT LL researchers and analyze their experience of such collaborations. Forty-one PIs have answered the survey. Forty-six percent of the respondents have been engaged in some form of collaboration with MIT LL, and the vast majority of them were satisfied with the collaboration. The vast majority of PIs (82.5%) expressed high enthusiasm about possible collaboration, but those who were not currently collaborating didn't know how to start the collaboration. In fact, 37% of the PIs didn't even know researchers in their area at MIT LL. In open response slot, multiple PIs expressed interests to give a talk or have a tour in the lab, but were not aware on how to engage. PIs who were not satisfied with MIT LL interactions mentioned mismatch of expectations as the main reason for the aborted collaborations. Other concerns include an unstable (interrupted mid-year) funding cycle and a large number of demos expected from the CSAIL group.

In the future, we propose that joint projects be selected via an open bidding process in the Laboratory. The selection committee, appointed by CSAIL and MIT LL leadership, has to announce ahead of time topics of interest and submission timeline. The projects have to be selected based on their mutual value to MIT LL and CSAIL, with special emphasis on their potential to lead to external funding. Additional criteria should focus on developing new collaborations and expanding the efforts to new research areas.

3.2 COLLABORATIONS TO ENHANCE CSAIL UNDERGRADUATE EDUCATION

For years, MIT LL has taken advantage of its strong ties to MIT campus to employ MIT IV-A undergraduate summer interns. Recently, MIT LL has become involved in the MIT UROP and UPOP initiatives, and this involvement has been steadily increasing. While employing undergraduates through these avenues has clear benefits to both MIT LL and the students, until recently there has been little explicit collaboration between CSAIL and MIT LL designed expressly to enhance undergraduate education for EECS students. A few notable exceptions include (1) the Computational Research in Boston and Beyond (CRIBB) seminar series, and (2) Capstone and IAP projects hosted by the MIT Lincoln Laboratory Beaver Works Center.

The CRIBB monthly seminar (<http://math.mit.edu/crib/>) was started in 2005 to promote the exchange of ideas among computational scientists at MIT and the broader community. The seminar co-organizers are Dr. Dreher (CSAIL & LNS), Prof. Edelman (Mathematics Dept. and CSAIL), Dr. Hill (EAPS), Prof. Johnson (Mathematics Dept. and RLE), Dr. Kepner (MIT LL, CSAIL, and Mathematics Dept.), and Dr. Reuther (MIT LL). Since 2005, CRIBB has hosted over 90 talks and has become a focal point for the computational science community at MIT. CRIBB has also provided a key venue for work performed at the Massachusetts Green High Performance Computing Center (MGHPCC.org). The seminars are conducted on campus at the Stata Center and attract several EECS undergraduates and graduate students.

The MIT Lincoln Laboratory Beaver Works Center is part of both the MIT School of Engineering and MIT LL. Although we assume most readers are familiar with the Beaver Works Center, Appendix B provides a short overview that describes the center, its goals, and its outreach to MIT undergraduates. One of the main innovations of the center is the use of project-based education to enhance the MIT undergraduate education experience.

Several dozen Beaver Works capstone projects have been run together with the Aero Astro, Mechanical Engineering, and other MIT departments. Capstone projects are one-year development projects executed in conjunction with a two-semester course, in which students work together to carry an idea from initial design all the way through to prototype fabrication and testing. To date, only one capstone project has involved CSAIL, namely, the Robust Communications for Autonomous Swarms project.

Professors who have participated in capstone projects have remarked that these projects have brought value to the students beyond what MIT typically offers in an undergraduate course. In particular, capstone projects provide the opportunity to engineer and build working prototypes that address real-world problems, and hence provide both scope and depth, thereby enriching the undergraduate education experience. Moreover, there is a shift from the professor serving as assessor to the professor serving as coach.

MIT LL's role has been to establish funding (for a course Research Assistant), access to real problems/requirements, expert mentors, and to make available the resources of the Beaver Works Center, which is located close to MIT campus at 300 Technology Square (NE45-202) in Kendall Square, Cambridge.

We recommend increased collaboration through tailored courses, MIT LL mentors, and Beaver Works capstone projects. CSAIL undergraduate students are very enthusiastic about taking the knowledge gained from their classes and applying it to practical problems. Through greater collaboration, students could be exposed to even more engineering challenges that they can learn how to address in lecture-type, lab-type, or project-oriented classes.

MIT LL mentors could provide more real-world context to the problems that students are tackling. MIT LL mentors could provide career guidance and exposure to the students. Furthermore, new courses and/or new content in old courses could be developed to strengthen the course curriculum in the areas of relevance to MIT LL missions.

3.3 COLLABORATIONS TO FACILITATE TECHNOLOGY TRANSFER FROM CSAIL TO MIT LINCOLN LABORATORY

CSAIL is a source of tremendous technology research and development. It has world-renowned academic researchers working in the areas of computer science and artificial intelligence, spanning fundamental research, innovative systems, and advanced technologies. Moreover, the outstanding work going on at CSAIL has never been more relevant to the many aspects of national security that MIT LL is chartered to address.

In general, greater awareness by MIT LL staff of the work being done on campus would help to prime this conduit, as would greater awareness on the part of CSAIL faculty of the national security problems being addressed at MIT LL. We feel, therefore, the single most effective measure to creating and sustaining a technology pipeline will be to embed MIT LL staff in CSAIL research projects that align with MIT LL interests. This will develop long-term, trustful collaborations with CSAIL, will help CSAIL better achieve its research goals, will allow MIT LL staff to become very familiar with CSAIL technologies and capabilities, will help attract top CSAIL students to MIT LL, and will allow for more effective situational awareness and cultural alignment between the two organizations. Funding for these staff would come from a combination of TO (initial seeding) and established programs.

In addition, increase of the scope of research collaborations between the laboratories and joint educational initiatives described above will necessarily increase technology transfer between the laboratories.

3.4 COLLABORATIONS TO CREATE FUNDED PROJECTS TO ADDRESS GRAND CHALLENGES IN CSAIL DOMAINS

Four study subcommittees were formed to explore potential “grand challenges” in key areas of mutual interest to CSAIL and MIT LL. The four areas examined were:

1. Autonomy-in-motion
2. Autonomy-at-rest
3. Advanced computing architectures
4. Cyber security

Autonomy-in-motion refers to robotics and machine autonomous platforms. Autonomy-at-rest refers to machine learning systems and algorithms, such as recommender systems, and “Big Data” analytics. Advanced computing architectures broadly refer to computing hardware and system software to address “Big Data” computing challenges. Cyber security includes cryptography more narrowly and cyber resilient systems in general. A few areas, such as communications, outside of these four domains were also considered.

The major goal of each subcommittee was to identify collaboration opportunities that could result in externally funded programs, and then to recommend next steps in pursuit of these opportunities. The other objective was to begin the collaborative dialogue between CSAIL and MIT LL, to become familiar with each other, and to share ideas and perspectives.

3.4.1 Autonomy-in-Motion (Robotics) Grand Challenges and Recommendations

Subcommittee Members

Prof Daniela Rus (EECS/CSAIL)
Prof Julie Shah (Aero Astro/CSAIL)
Dr. Beijia Zhang (MIT LL)
Dr. Michael Boulet (MIT LL)
Mr. Ken Gregson (MIT LL)

Overview

The DoD distinguishes autonomy-in-motion (including robotics and mobile autonomous systems) from autonomy-at-rest (including Big Data analytics, human-machine interfaces, etc.) Although the areas are clearly complementary, this study addresses them separately and also identifies beneficial or synergistic connections between them. Robotics is an extremely active research area rich with opportunities to both advance the state of the science and demonstrate new operational capabilities. DARPA has used Challenge Events (2004/2005 Grand Challenge, 2007 Urban Challenge, 2012–15 DRC)

to focus and accelerate development of autonomous vehicle navigation and humanoid robotics. As a result of these high profile events and increasing commercial interest in automated transportation and delivery systems, significant advances have been achieved in autonomous locomotion, navigation, perception, planning, and interaction. However, even the best systems today do not operate with the speed, robustness, endurance, scalability, and trust needed for widespread adoption and use. Fewer yet can operate in the complex, uncertain, often geographically distant, and hostile environments required for many DoD missions. Therefore, we propose establishing a significant joint CSAIL+MIT LL collaboration for research and demonstration of Autonomous Operation in Difficult Environments.

Nature of the CSAIL-MIT LL Collaboration

A CSAIL-MIT LL collaboration would ideally leverage the particular strengths of both laboratories to strategically span the full technology spectrum from basic science to systems-level integration and demonstration, in a representative mission context as represented in Figure 3 below. Since its inception, CSAIL has been a world leader in artificial intelligence and robotic systems research. Likewise, MIT LL has been at the forefront of integrated systems and advanced sensor development for the DoD. Both institutions have expertise that covers the entire gamut of technologies utilized in autonomous systems.




Figure 3. Combining strengths of both institutions.

Another intricacy of the CSAIL-MIT LL relationship is classification, review, and releasability of research performed under DoD sponsorship. We propose addressing a challenge that maximizes opportunities for publication and to structure the activities in such a way as to minimize such restrictions on the CSAIL portion of any joint projects.

Proposed Challenge Problem: Extended Reach for Continuous Operations

With the foregoing in mind, we identify Continuous Resupply for Humanitarian Assistance and Disaster Relief as a crucial and suitable challenge. This concept is depicted notionally in Figure 4 below. At the highest level, this can be framed as a logistics problem addressed by a distributed, layered, and linked network of heterogeneous autonomous systems.




Figure 4. Continuous resupply for humanitarian assistance and disaster relief.

It provides both a representative difficult environment and an accessible venue to concentrate on developing and demonstrating the critical elements of speed, robustness, endurance, scalability, and trust that should be the focus of this work. In addition to HADR, this general problem has analogs in many

commercial and military applications and missions that serve to further broaden its interest. It also has the benefit of presenting many subproblems that may be independently addressed and later integrated into larger systems demonstrations. Three example subproblems are outlined in more detail in the following sections for Standoff System Architecture, Forward Deployed Ground Support, and Legged Autonomous Ground Vehicle.

Subproblem #1: Autonomous Standoff System Architecture

Rapid and effective response is critical in humanitarian and disaster relief events, and it has been repeatedly demonstrated that the world's current framework and systems cannot react as efficiently as desired. Sufficient infrastructure is often lacking in disaster areas or devastated by the event itself. A rapidly deployable distributed, autonomous standoff system architecture that could alleviate those issues and scale as needed is shown notionally in Figure 5 below.




Figure 5. Autonomous standoff system architecture.

Subproblem #2: Forward Deployed Ground Support

A forward deployed ground system would provide the ability to extend the reach of the architecture further inland. This system could have unique capabilities and serve as a processing, exploitation, and command and control node. In addition, it could also provide needed infrastructure support, such as a

base station, to establish communications in both directions and store/forward needed information in the event of communications disruptions to support disconnected operations.

The ground system would function in similar fashion to the maritime standoff rigid hull inflatable boat (RHIB) as a staging depot for resupply of the forward-acting elements of the mission. A chain of such ground systems could be emplaced rapidly to provide broad area coverage based on determined need and demand.

Subproblem #3: Legged Autonomous Ground Vehicle (Jaguar)

Many autonomous unmanned platforms already exist, but in the ground domain, legged robots have substantial advantages for navigating through complex terrain and environments. There are many situations where wheeled ground vehicles and unmanned aerial vehicles (UAVs) encounter challenges that prevent their use or reduce their effectiveness. An agile legged platform that can autonomously navigate at high speed with long endurance/power efficiency, quiet operation, and reasonable payload capacity would fill a significant gap and could play an important role in both civilian and military applications. A proposed system based on the success of the MIT Cheetah is notionally shown in Figure 6 below.




Figure 6. Legged autonomous ground vehicle (Jaguar).

Canonical Autonomy Architectures

Embedded within this challenge problem are two orthogonal frameworks or architectures for autonomous systems. Scalability may be achieved through a layered autonomy architecture that greatly expands the scope of human control through the use of high-level mission tasking and distributed autonomous execution. The limits of endurance for individual systems are likewise overcome through a linked network or chain of autonomous systems to extend the effective reach of the overall system. These architectures are illustrated notionally in Figure 7 below.




Figure 7. Two orthogonal frameworks for autonomous systems.

Next Steps

One hurdle repeatedly encountered in preparing joint proposals for sponsored research is the short timeline given for proposal submission. Potential options to address this challenge include establishing a standing quick reaction proposal team that can rapidly turn around proposals. An alternative, or perhaps supplementary, approach would be to have proposals on the shelf prepared in advance and ready to submit or tailor as opportunities arise. We recommend that both approaches be pursued and a proposal team begin the work now of decomposing the selected challenge problem into subproblems suitable for joint development and demonstrations that we can offer to potential sponsors for funding in white paper form independently of BAAs (Broad Area Announcements) or other calls for proposal. This also helps

address the difficulty for MIT LL to respond to BAAs and increases the likelihood that a joint CSAIL-MIT LL proposal could be funded.

Links to Other Subcommittee Domains

We earlier observed that the areas of autonomy-in-motion and autonomy-at-rest are highly complementary. While they are being addressed independently in this study, we would like to point out some connections between them that could be exploited for mutual benefit. Scalable autonomous systems such as those imagined here clearly need an advanced human machine interface in order to be employed effectively and, ultimately, trusted. Mission-level tasking must also be communicated to the distributed autonomous agents creating and conducting lower-level execution plans. Those plans could be provided directly by human controllers or through autonomous agents running streaming analytics on big data. Of course, cyber security is a factor that needs considering throughout any such system in order to be resilient and trusted.

What Success Might Look Like

From a financial perspective, an initial start could easily translate into \$3–5M of CSAIL-MIT LL joint funded projects annually over the next 3–5 years in the area of autonomy-in-motion alone.

Another measure of success would be the extent of the collaboration both at CSAIL and MIT LL. A recent survey of campus research activities revealed 31 PIs are conducting robotics research. Many were unaware of the number or scope of the projects currently being pursued. A similar situation exists at MIT LL, where 29 proposals were submitted for internal funding to the Autonomous Systems Line with more proposals per unit of available funding than other investment areas. Both situations are indications of the combined high level of interest and excitement about robotics and how diffuse that interest is across Departments/Laboratories and Divisions/Groups at our respective institutions. An inspiring challenge problem with many and varied opportunities to participate could better cohere research interests at each institution and between them as well.

3.4.2 Autonomy-at-Rest (Analytics) Grand Challenges and Recommendations

Subcommittee Members

Robert Bond (MIT LL)
Regina Barzilay (MIT)
Sanjeev Mohindra (MIT LL)
Charlie Dagli (MIT LL)
Michael Yee (MIT LL)
Mike Hurley (MIT LL)
Arjun Majumdar (MIT LL)
Fred Waugh (MIT LL)

Overview

The Autonomy-at-Rest subcommittee explored opportunities for technical exchange, collaboration, and synergy in the broad areas of machine learning, advanced analytics, man-machine interfaces, and computational architectures. In general, the committee found that there are numerous areas of mutual interest and potential collaboration. The committee identified grand challenges in both early indications and warnings (I&W) for missions such as counter-weapons of mass destruction (C-WMD) and time-critical decisions (TCD) for missions requiring rapid deterrence or interdiction.

The autonomy-at-rest research thrust tackles the challenges posed by the complexity of data and the need to exploit complex data for decision-making under uncertainty. The application areas of interest to MIT LL researchers span the full “kill chain” in data exploitation, decision-making, and command and control for national security applications. The challenges are often characterized by the lack of ground truth and associated labeled data. Thus, the autonomous techniques developed under this research thrust must address the twin problems of the complexity of the data space and the scarcity of truth data. These two problems make autonomy-at-rest an exciting area for ground-breaking research in “low-resourced” machine learning, computational architectures, and databases.

A canonical autonomous architecture for exploiting complex data is shown in Figure 8.




Figure 8. Autonomous architecture components.

The key elements of the architecture and potential collaborative research and development areas are

1. Smart Sensing

The smart-sensing research area encompasses development of sensor networks that can be tasked to cooperatively sense the environment. We envision both physical sensors to sense the real world and virtual sensors to sense the web, the deep web, and the dark web. Two important areas of research are autonomous sensing networks and putting rudimentary perception capability at the sensor to improve sensing, reduce power consumption, and reduce the data-flow requirements.

2. Autonomous Analytics at Scale

Autonomous analytics at scale are needed to operate on both real-time data streams and data at rest. Interesting research areas include making advances in image, text, audio, and video processing; graph analytics; machine learning in low resource situations with missing or mislabeled data; multimodal data fusion and correlation, pattern recognition, patterns-of-life, and anomaly detection. For real-time analysis, online versions and approximate algorithms need to be researched and matured.

Several faculty and research scientists have research agendas that overlap this thrust.

3. Multimodal, multisource knowledge bases

This is already a strong area of research at CSAIL. Research is needed on how to store and efficiently query multimodal, multisource data. Areas of interest include in-memory databases, embedded database computing, distributed databases, and graph databases. In the area of graph databases, research is needed on probabilistic pattern matching on graphs, as well as development of autonomous agents that reason over the stored data.

4. Virtual Aide

Data-driven decision-making presents a rich area of research encompassing computational workflows, decision processes, data visualization, and the human-machine interface. Particularly interesting research areas are big data and big graph visualization, visualization to aid anomaly/outlier detection, and provenance management & explanation systems.

This is another strong area of research for CSAIL.

5. Autonomic Architectures

Today's systems are complex and present a management challenge. Building on CSAIL's strength, research in this area can be focused on resource allocation and scheduling algorithms to feed self-management and self-optimization; semantic API and workflow planning to compose services into workflows; and autonomic cyber security from the ground up. Current CSAIL research into next generation cluster managers such as Mesos can be leveraged and expanded for developing autonomic architecture components.

Next Steps

A next step would be to discuss specific parts of the problem among subject matter experts at CSAIL and MIT LL. Table 1 provides a partial list of candidate principal investigators (PIs) and potential research collaborations.

TABLE 1
Partial List of Principal Investigators and Potential Research Collaborations

Area	Potential CSAIL PIs	Potential MIT LL PIs	Potential Research Collaborations
Text Analytics and Machine Learning	<ul style="list-style-type: none"> • Regina Barzilay • Tommi Jaakkola • Stephanie Jagelka • Tamara Broderick • Leslie Kaelbling 	<ul style="list-style-type: none"> • Olga Simek • Danelle Shah • Charlie Dagli 	<ul style="list-style-type: none"> • Develop methods to detect population vulnerable to recruitment/bullying based on their social media footprint • Develop indicators for recruitment/bullying activity on social media
Computer Vision and Machine Learning	<ul style="list-style-type: none"> • Bill Freeman • Antonio Torralba • John Fisher • Tommi Jaakkola • Stephanie Jagelka • Tamara Broderick • Leslie Kaelbling 	<ul style="list-style-type: none"> • Arjun Majumdar • Mike Hurley • Ben Smith 	<ul style="list-style-type: none"> • Generalized object detection/classification and activity monitoring in commercial satellite imagery • Distributed imaging and scene understanding using UAV swarms <ul style="list-style-type: none"> – Processing of non-contiguous field-of-view sensor data – Optimization of sensor tasking between multiple platforms • Video summarization • The use of deep learning and other methods to 3D object recognition in 3D lidar and EO/IR point-clouds and imagery • Video and multimedia-based search algorithms that operate in low-resourced and noncooperative (untagged) domains
Knowledge Representation and	<ul style="list-style-type: none"> • V. Manasinghka • Sam Madden 	<ul style="list-style-type: none"> • Fred Waugh • Sanjeev Mohindra 	<ul style="list-style-type: none"> • Probabilistic knowledge representation and

Area	Potential CSAIL PIs	Potential MIT LL PIs	Potential Research Collaborations
Databases		<ul style="list-style-type: none"> • Jeremy Kepner 	<ul style="list-style-type: none"> probabilistic computing • Storing and efficiently querying multimodal data • System for autonomously generating insight from data based on initial human input • Computing on masked data; privacy preserving data-mining • Data visualization • Provenance tracking and explanation services • Develop and implement array-based API and operations on SQL, NoSQL data stores and data streams
Architecture	<ul style="list-style-type: none"> • Matei Zaharia 	<ul style="list-style-type: none"> • Michael Yee • Jeffrey Hughes • Sanjeev Mohindra • Jeremy Kepner 	<ul style="list-style-type: none"> • Resilient software services that adapt to failure/changes in the hardware infrastructure • Schedulers that allow self-management, optimization, and repair of resources • Dynamic task and data parallel mapping of hierarchical data arrays and data streams

3.4.3 Advanced Computing Architectures Grand Challenges and Recommendations

Subcommittee Members

Prof. Arvind (CSAIL)
Dr. Vijay Gadepally (MIT LL)
Dr. Jeremy Kepner (MIT LL, co-chair)
Dr. Huy Nguyen (MIT LL, co-chair)
Dr. Albert Reuther (MIT LL)

Additional Input from

Prof. Edelman (CSAIL), Prof. Leiserson (CSAIL), Prof. Madden (CSAIL), Prof. Stonebraker (CSAIL), and Dr. Joseph Campbell (MIT LL).

Overview

Computer architecture and advanced computing have long been areas of collaboration between CSAIL and MIT LL and have experienced a significant increase in activity catalyzed by the Beaver Works Center, the Massachusetts Green High Performance Computing Center, the Center for Engaging Supercomputing, MIT LL Line investments, and external sponsor investments. Interest in the field overall has increased because of the vast volume, velocity, and variety of Big Data that now exists and is expected to further expand as the Internet-of-Things era takes hold. CSAIL and MIT LL independently do a large amount of research in these areas, and thus collaborations are quite natural.

Volume

The need to store and process ever-increasing volumes of data had been the driving force behind the MIT SuperCloud research effort. Originally developed by MIT LL as the LLSuperCloud, this technology is the engine for a number of collaborations. More specifically, the MIT SuperCloud is the joint software environment used by both MIT LL and MIT to run their supercomputers. Further development and transition of this technology has been funded externally as a joint collaboration between CSAIL and MIT LL. Finally, the MIT SuperCloud is used to support a wide range of collaboration across MIT LL and campus.




Figure 9. MIT SuperCloud software stack used across MIT Lincoln Laboratory and MIT supercomputers.

Velocity

Processing data at ever-increasing rates requires new architectures that can efficiently stream data. The MIT LL Line-funded CSAIL MIT LL Director's Initiative collaboration in low-power embedded analytics seeks to develop and utilize new computing architectures to accelerate big data analytics and enable low-power in-situ sensor processing with novel in-storage computing architecture. The approach is

to refactor nonvolatile data storage, networking, and processing functions into field-programmable gate arrays (FPGAs) to improve throughput and power efficiency.




Figure 10. New architecture for low-power off-grid near-sensor processing.

Variety

The enormous variety of data is the hardest challenge in developing modern data processing systems. Several joint CSAIL-MIT LL collaborations are taking on this challenge. The centerpiece of these efforts is the BigDAWG database federation technology that aims to allow highly distinct database to interoperate.




Figure 11. BigDAWG federated database architecture.

Identified Areas for Collaboration

CSAIL's ability to generate innovative ideas has led to a wide range of sponsor interest in the above technologies. This has led to a number of opportunities to conduct additional collaborative research.

Next Steps

We should continue to leverage strengths of each organization. CSAIL has world-renowned faculty, brilliant students and post-docs, and individuality. MIT LL has extensive expertise in systems analysis and design, strong interdisciplinary collaborations, large resources/infrastructure, access to hard/important problems, and expertise with handling of diverse information. CSAIL faculty and MIT LL staff should work to build computing capability teams that enhance their research vision and enables undergraduates and masters' students to meaningfully contribute to leading-edge research. MIT LL staff should be encouraged to spend significant time at CSAIL.

3.4.4 Cyber Security Grand Challenges and Recommendations

Subcommittee Members

Prof. Srinivas Devedas (EECS)
Dr. Howard Shrobe (EECS)
Dr. Rob Cunningham (MIT LL)
Dr. Hamed Okhravi (MIT LL)
Dr. Michael Vai (MIT LL)
Dr. David Whelihan (MIT LL)

Overview

Department of Defense (DoD) systems are increasingly the targets of deliberate and sophisticated cyber-attacks. To assure mission success, military systems must be entrusted to perform their intended functions, prevent attacks, and operate while under attack. The DoD has thus directed that cyber security must be integrated into the full military system life cycle.

Over the last few years, MIT LL's Cyber Security and Information Sciences Division has been developing cyber security technologies that span from modeling and simulation to hardware devices. Funding sources for these developments vary from internal (e.g., line) to external (e.g., DARPA) funding.

Compared to many other MIT LL mission areas, cyber security is still in its infancy. As such, many problems warrant research and development activities that span across 6.1 (basic research) and 6.2 (applied research) areas. Figure 12 summarizes the subcommittee's observations in the evolution of secure embedded computing. We have witnessed that many systems are developed by contractors and commercial developers into "one-of-a-kind" secure systems.

Judging from our involvement with leading-edge research programs (e.g., the DARPA High-Assurance Cyber Military Systems), the future of cyber security lies in open and universal technologies. Currently, these technologies, such as formal verification of protocols and software designs, are dominated by academic research institutes (e.g., CSAIL). This situation creates a perfect opportunity for MIT LL and CSAIL to collaborate, as they naturally contribute to different aspects of solutions to cyber security challenges.




Figure 12. Secure embedded computing evolution: past, present, and future. Academia dominates the future open and universal security technology development.

Table 2 summarizes the current collaborations between CSAIL and the MIT LL Cyber Security and Information Sciences Division. The highlighted projects are directly relevant to cyber security. The rest of the projects, while not directly related, are also often applicable to certain aspects of cyber security.

TABLE 2
Ongoing Collaboration between CSAIL and the Cyber Security and Information Sciences Division

Topic	Lincoln POCs	CSAIL POCs
Next-Gen Speaker and Language Recognition Algorithms	Weinstein/Campbell	Glass/Dehak
Julia: Big Data Extensions and Analytics	Gadepally/Campbell	Edelman
Image Analysis	Miller	Polina Golland
Automated Corpora Generation for Software Analysis*	Leek	Zeldovich
Memory Corruption Attacks and Defenses*	Okhravi	Shrobe/Rinard
Cyber Intrusion Analysis*	Carter	Shah
Functional Encryption*	Yerukhimovich	Goldwasser
Secure Hardware*	Whelihan/Vai	Devadas
Advanced Database Technologies	Gadepally	Stonebraker/Madden
MIT SuperCloud benchmarking	Kepner	Leiserson

*Cyber security-related projects

MIT LL has been developing security building blocks, particularly those relevant to the security of embedded computing, which is a key area of military systems. Results from these efforts, such as LOCKMA (Lincoln Open Cryptographic Key Management Architecture), PUF (physical unclonable function), and S-COP (security coprocessor) have been demonstrated and transferred to government programs (ONR, AFRL, Draper, etc.). We describe below a collaboration example in this area of secure hardware, which is the development of a JIT-ST (Just-In-Time Secure Thread) processor.

As shown in Figure 13, the JIT-ST processor is a novel key-centric processor architecture in which each piece of data or code can be protected by encryption while at rest, in transit, and in use. Using embedded key management for cryptographic key handling, the processor permits mutually distrusting software written by different entities to work closely together without divulging algorithmic parameters or secret program data. Since the architecture performs encryption, decryption, and key management deeply within the processor hardware, the attack surface is minimized without significant impact on performance or ease of use.




Figure 13. Just-In-Time Secure Thread (JIT-ST) processor performs encryption, decryption, and key management deeply within the processor hardware to enable mutually distrusting software to work closely together without divulging secret data. Attack surface is minimized without significant impact on performance or ease of use.

The JIT-ST team consists of MIT LL staff members and CSAIL researchers in Professor Srinivas Devadas' Computation Structures Group. The objective of the collaboration is to jointly expand knowledge in the area of Trusted and Secure Computing. The team has been working to understand the ramifications of technologies such as silicon PUF (physical unclonable functions), just-in-time encryption/decryption, path oblivious RAM (ORAM), and advanced key management on overall system security and usability/performance.

Specific accomplishments achieved in this collaboration are summarized as follows: MIT LL's secure computing knowledge has been expanded through a collaborative participation in weekly meetings. The insight gained has been used to guide specific technology developments of the JIT-ST processor. When appropriate, MIT LL has used general use case and threat model discussions to bring industry and DoD perspectives to the technologies being developed by Professor Devadas' group.

Another successful collaboration example is the cyber security evaluation and prototyping project. A team of CSAIL faculty, researchers, students, and MIT LL staff collaboratively performs cyber security research at the Beaver Works Center. The research challenges fundamental hypotheses behind current defenses, and then designs and implements new defenses. Results have been published in prominent security venues such as ACM CCS'14 and IEEE S&P (Oakland'15).

Identified Areas for Collaboration

Cyber security is of strong interest and importance to MIT LL, both near term and long term. As mentioned above, MIT LL and CSAIL have complementary strengths in basic research and system-level applied research and development, both of which are essential for both near- and long-term projects that are of importance to national security.

The subcommittee has identified a few ideas for moving forward in our collaboration. The first idea is to further enhance the benefits of current interactions. We propose to expand the two-way conduit between basic (6.1) and applied (6.2) research by opening more CSAIL student positions (e.g., RAs) on MIT LL projects (at Lexington and/or the Beaver Works Center). These projects could lead into thesis research projects, which are co-supervised by CSAIL faculty and MIT LL staff.

Another idea is to establish short-term (e.g., 3-month terms) MIT LL staff sabbaticals to be fully embedded into the CSAIL research environment. Currently, this type of opportunity exists for CSAIL faculty and/or researchers, but is unavailable for MIT LL staff.

Last but not least, we recommend the establishment of an open channel for communications between CSAIL and MIT LL. Currently, clear guidance on information sharing does not exist, to our best knowledge. For example, an overarching release process for MIT LL staff to share information with CSAIL researchers would be extremely helpful for discussions.

The subcommittee has also decided to pursue, as a unified team, extensive collaborative R&D projects sponsored by government/DoD agencies. The vision is that both CSAIL and MIT LL contribute different aspects of a big picture. We can build on the current CSAIL Director's Initiative, leverage technology building blocks, and target the development of system-level secure and resilient technologies.

The subcommittee has come up with a concept on the development of secure application-specific systems (SASS). This is motivated by the observations illustrated in Figure 14. The DoD has directed that cyber security must be integrated into system life cycles, as it has proved ineffective to secure a system using a post-design or post-implementation approach. Currently, developers routinely perform size, weight, and power (SWaP) and throughput tradeoffs to meet specific mission needs, but leave security to be treated as an "add-on." As a result, system developers either attempt to receive a waiver on security or rely on proprietary and outdated secure processors to protect critical information and technologies.

The design of security for an embedded system is challenging: security requirements are rarely accurately identified at the start of the design process, so engineers tend to focus on well-understood functional capabilities rather than stringent security requirements. In addition, they must provide security with minimal impacts on the system's SWaP, usability, cost, and development schedule.

The subcommittee proposes the vision of developing a design methodology and prototype components for composable, secure embedded processors, with the goals of ensuring security while providing long-term extensibility and mission adaptability. This methodology, which we refer to as SASS, consists of application-driven security requirements and implementation and language-level HW/SW co-design for functionality and security. The objective is to enable developers to work in a design space that not only provides the required throughput and SWaP dimensions of embed systems, but also security and anti-tamper attributes.




Figure 14. Secure application-specific systems' (SASS) vision; the mission system design space will be expanded to include a dimension in security. The objective is to enable the co-design of functionality and security from the get-go.

This project will provide open source SASS technologies that are formally verified and publicly vetted. The AES (advanced encryption standard) encryption algorithm is a great example of publicly vesting benefits; its vulnerabilities are well-known so that proper measures can be taken to defend against threats.

Figure 15 shows the SASS design flow that develops and builds application-specific secure systems a la carte. This security system hardware-software co-design environment will consist of threat models and concept of operation inputs, compilers/synthesizers, toolboxes, and IP (intellectual property) libraries. The output of the toolchain targets platforms (custom or COTS) that provide the proper functionality and security to support mission objectives.




Figure 15. SASS program concept. A tool set and libraries will be developed so that system developers can determine security requirements by analyzing applications and potential threats. The security requirements become the inputs of the system, which performs automatic or semi-automatic SE/HW security co-design. The result is synthesized (HW) and compiled (SW) into application-specific secure system architecture.

Next Steps

1. Create a plan to engage more CSAIL students on MIT LL projects.
2. Investigate and propose a plan to enable short-term MIT LL staff sabbaticals at CSAIL.
3. Investigate and propose a plan to establish open channels for communications between CSAIL and MIT LL.
4. Develop SASS development tasks, and identify who would participate/lead each of the tasks.
5. Develop a SASS execution plan.
6. Explore external funding opportunities.
7. Convene a joint team to contact potential funding sources.

3.5 COLLABORATIONS TO CREATE ON-SITE RESEARCH AND DEVELOPMENT OPPORTUNITIES BETWEEN MIT LINCOLN LABORATORY AND CSAIL

A few of the study members noted that CSAIL faculty would be receptive to having more assistance from MIT LL staff to oversee undergraduate and master's-level research and effectively integrate it into their research agendas. Following up on this, the LLSC (Kepner) proposed the concept of a "Lincoln Computation Institute" whereby MIT LL staff would spend 50% of one year working with a student on campus to advance a computational capability relevant to an MIT LL mission. The project would adapt and scale a computational capability on a campus LLSC system and then transition that capability to a MIT LL LLSC system so it could be used to support MIT LL sponsored research. Setting this up as a rolling application process would allow MIT LL and CSAIL to align the timing to when students are available (this is often a narrow window). LLSC would vet the computational aspects of the project while the associated Line or Division representatives would assess the technology and mission relevance.

4. SUMMARY

The committee is confident that strong research and development collaborations between the laboratories will greatly benefit both MIT LL and CSAIL. The committee found that existing collaborative projects served as valuable venues for transitioning technology and expertise from CSAIL to MIT LL while at the same time providing CSAIL with insight into real-world national security problems. After analyzing the scope of current collaborations and significant interests in both laboratories, the committee recommends to develop new engagement strategies aiming to expand the connections between the two laboratories. Such measures include open bidding process for collaborative projects, mechanisms for creating new collaborations, efforts towards joint externally funded projects, expansion of CSAIL involvement in Beaver Works, and long-term and/or sustained periodic visits of MIT LL researchers on campus.

This page intentionally left blank.

APPENDIX A

CSAIL DIRECTOR'S INITIATIVE

The CSAIL Director's Initiative emerged from a discussion in 2013 between the Director of MIT Lincoln Laboratory, Dr. Eric Evans, and the Director of CSAIL at that time, Prof. Anant Agarwal. The idea was to create a longer term research venue that would allow CSAIL professors to undertake sustained research and development projects in areas of mutual interest between CSAIL and MIT LL. Each project was designed to be a collaboration, with both CSAIL and MIT LL PIs. The nominal project duration was set at three years, and three broad research areas were established:

1. Low-Power Electronics
2. Machine Autonomous Systems
3. Cyber Security, including (a) cryptography and (b) secure hardware systems

LOW-POWER ELECTRONICS: LOW-POWER EMBEDDED ANALYTICS

The goal of this campus collaboration is to explore and demonstrate a new “near-data computation” architecture that could lead to 10× boost in performance and power efficiency for many big data and applications. Near-data processing is about migrating selected processing operations from the central processing unit (CPU) of traditional computer into the transfer paths with storage / network, where they can be implemented with application-specific logic blocks to vastly improve performance. For applications with strict requirements on SWaP, such as remote intelligent sensing and airborne computer-on-watch, this enables new operations not possible before, while for conventional office settings, it enables economical solutions for scaling up performance.

In the first year, the MIT LL team performed system analysis and developed application concepts. In parallel, professor Arvind’s research group started developing the BlueDBM system, where DBM stands for database management system. A BlueDBM node consists of two custom-built flash/communication modules interfaced to a commercial FPGA board that plugs into a server computer. The FPGA performs three functions: (1) accelerate selected CPU tasks, (2) manage flash storage, and (3) provide a high-speed custom network to other BlueDBM nodes. The development was co-sponsored by Quanta Research, Samsung, Intel, and Xilinx. In the second year, preliminary measurements on BlueDBM revealed storage bandwidth of 2.4 GBytes/s for each 1 TB node, and communication links of 80 Gbits/s per node, exceeding commodity commercial capability by 5× and 8×, respectively. We also developed and benchmarked simple application kernels to better understand design tradeoffs on this new architecture. Two publications were submitted on BlueDBM.

This year (2016), we plan to further develop infrastructure firmware and software for the system. Additionally, at least two reasonable-sized applications will be developed to demonstrate benefits of near-

data computation. One application will involve interfacing with D4M, a MIT LL technology for big data algorithm development from within a MATLAB environment. The other application will be related to SWaP-constrained computer vision and machine learning. We'll also pursue a few opportunities for using BlueDBM as a proof-of-concept development environment for near-data processing solutions to MIT LL and external programs.

MACHINE AUTONOMOUS SYSTEMS

DDAGOUS – Data-Driven Autonomy for Group Operations in Uncertain Scenarios – is a CSAIL-MIT LL collaboration in the area of autonomy, with Professors Daniela Rus and Julie Shaw as CSAIL principal investigators. The effort seeks breakthroughs in autonomy to effectively utilize DoD unmanned platforms with reduced human intervention. The research bridges the currently disparate fields of machine learning, control theory, information theory, online discrete algorithms, and human-machine systems to develop new algorithms that can deliver real-time performance for autonomous DoD systems. The program's three research thrusts are (1) dynamic decision-making and coordination of humans and autonomous agents under communication and information uncertainty, (2) persistent autonomy in dynamic environments with online data, and (3) using historical data in real-time decision making: optimized semantic data compression.

The first-year (FY14) effort focused on continuing corset research in Professor Rus' research group. Coresets are scalable algorithms for segmentation and event summarization of streaming high-dimensional data, such as geospatial position or video. The work produced two papers: Coresets for Visual Summarization with Applications to Loop Closure and Coresets for k-Segmentation of Streaming Data. Nicholas Armstrong-Crews, the MIT LL staff supporting the program, started work on applying the coreset video summarization algorithm to the problem of structure from motion, i.e., constructing three-dimensional point clouds from a translating camera sequence of two-dimensional images. However, briefings to the Autonomous Systems (AS) Line external review panel and the AS Line Stakeholder Panel failed to present clear research objectives, milestones, and technical plan. The proposed FY15 effort was ranked 10th by the Stakeholder Panel.

The second-year (FY15) effort incorporated Professor Shaw's research group in the collaboration. Additionally, Michael Park started to support the effort from MIT LL after Nicholas Armstrong-Crew's departure from the Lab. The FY15 AS Line external review feedback again noted that a clear plan, milestones, and progress against those milestones were not presented and that the two research groups could benefit from better coordination. In the months following the external review, Michael Park worked with the CSAIL collaborators to define a multi-UAV surveillance demonstration that incorporates both groups' research. Additionally, Michael Park made progress leveraging coresets for structure from motion using an airborne video dataset, although work remains to quantify the advantages over existing methods.

A continuing challenge for this effort has been presenting a focused research objective and plan to MIT LL and external reviewers. The FY16 research portfolio addresses this issue. In FY16, Mark Donahue became the MIT LL PI. The goals for the third year (FY16) are to merge vision and tactile

feedback to interact with UAVs, and then to conduct human-machine interface experiments to assess the effectiveness or hindrance of modality switching. The project will use a gaming environment for quick-turn, human-in-loop evaluation, and will explore using coresets with metrics targeted for human-machine interface applications. The project will culminate with a demonstration relevant to the Department of Defense, showing the effectiveness of merged vision and tactile feedback to interact with UAVs.

CYBER SECURITY: CRYPTOGRAPHY USING FUNCTIONAL ENCRYPTION

CSAIL-MIT LL collaboration with Professor Shafi Goldwasser and Arkady Yerukhimovich

Functional encryption is a powerful, recently proposed cryptographic primitive that combines access control with computation on encrypted data, allowing data owners to control exactly who can compute what based on their data. As such, functional encryption has great potential to address many of the security concerns that arise in today's outsourced computation environments, such as in cloud computing. Functional encryption can allow these data owners to allow certain processing to be done on their data while maintaining complete control over their data. The functional encryption research project aims to understand the current state of this powerful cryptographic tool, to help transition it closer to practice and to establish MIT LL as an expert in designing and developing schemes based on functional encryption, especially for DoD applications. Towards this goal, we are developing new models, security definitions, and constructions in order to find optimal tradeoffs between generality, security, and efficiency. We are working together with Professor Shafi Goldwasser at CSAIL and her students to mature this research, develop prototypes to demonstrate feasibility, and publish our results at appropriate academic conferences.

In the last two years of this project, we have built up our understanding of the current state of the art in functional encryption and identified avenues for advancing this research. In particular, we performed a survey of existing literature to understand the limitations of current approaches and have proposed several ways to circumvent the known bounds on current constructions. The main approach we have taken is to investigate nonstandard definitions for functional encryption to enable more efficient constructions. For example, in FY15 we showed a new construction of bounded-collusion attribute-based encryption, a special case of functional encryption focused on access control that is more efficient than any previously known construction.

In the third year of this program, we plan to further develop these techniques to build new functional encryption protocols. Additionally, we plan to prototype the constructed schemes as appropriate to demonstrate feasibility of their practical use. We plan to pursue research along the following directions. First, we will continue our work on functional encryption with relaxed security models. As demonstrated by our work on attribute-based encryption, relaxed security models, such as the bounded-collusion model, often allow for constructions using much simpler and more efficient building blocks. We aim to improve upon this work to build more efficient and practical functional encryption schemes. Second, we will investigate how to strengthen and improve functional encryption for specific functionalities of interest. In particular, we will focus on improving constructions of functional encryption

for inner product functionalities as this is a class of functions with many uses, and the several existing constructions do not achieve all the desired properties. Finally, we will consider how we can leverage nontraditional models for functional encryption to overcome the known limitations. Here we plan to consider such relaxations as symmetric-key functional encryption and interactive functional encryption to develop tools to address problems that cannot be addressed by standard functional encryption.

CYBER SECURITY: SECURE HARDWARE SYSTEMS

The advent of fully homomorphic computing (FHC) has changed our thinking about the limits of system security. Unfortunately, the extremely high computational overhead of FHC techniques renders them unusable in the vast majority of applications. However, getting as close as possible to achieving the chief property of FHC—that it never requires code or data to exist in plaintext—is still a worthwhile, if lofty, goal. The primary focus of the Just-In-Time Secure Thread Processor (JIT-ST) program has been to architect, design, and ultimately build a processing system comprised of a wide range of technologies that will keep data in plaintext for the shortest time possible, while still providing high performance. While it is currently technically infeasible to implement a FHC processor, which has a zero attack surface, we want to minimize the attack surface of the processor being designed without significantly impacting system performance.

Over the last two years, the MIT LL team has had extensive interaction with Prof. Srini Devadas and his group of graduate students as collaborators. By attending weekly meetings, effectively “embedding” with the campus team, the MIT LL team aided in their design of a secure Path-ORAM test chip and exchanged a lot of information, insight, and knowledge on a variety of topics, including advanced secure computer architectures such as Intel SGX, Physical Unclonable Function (PUF) enabled generation of stable, repeatable keys to enable unique hardware IDs, and uses for other MIT LL technologies such as cryogenic logics (a study funded under the ICE line). These are all leading edge (albeit academic) technologies that were extremely beneficial to our in-house effort: develop a refined JIT-ST processor architecture that could be deployed for a wide range of applications in various environments, which uses just-in-time encryption to bring strong encryption deep into the execution pipeline and minimizes attack surface.

Over the next year, we plan to continue our research collaboration with Prof. Devadas’ research group. In particular, we will pursue a few avenues to deepen our collaboration, including

- Sharing with Prof. Devadas’ team, the MIT LL-developed reference processor architecture based on the Sparc v8 processor for reviewing and collaboration purposes. In the near future, we intend to open-source this reference processor architecture for the broader research communities. This architecture is very suitable for the deep insertion of security technologies. This is within the scope of our proposed FY16 JIT-ST effort.
- Working with the Advanced Technology Division (Division 8) to obtain “ground-truth” on ring-oscillator-based PUFs by embedding relevant test structures in Microelectronics

Laboratory (ML) test runs. The resultant structures would be used to test some of Prof. Devadas' latest technologies to derive stable keys from device-specific characteristics. This project will benefit both Prof. Devadas and MIT LL. This project needs further discussion with Division 8.

This page intentionally left blank.

APPENDIX B

MIT BEAVER WORKS CAPSTONE PROJECTS

As described on the MIT Lincoln Laboratory Beaver Works Center (Lincoln Beaver Works) home page (<https://beaverworks.ll.mit.edu/CMS/bw/facility>), the Beaver Works Center conducts research and educational programs that strengthen and expand collaborative efforts between MIT LL and MIT campus. This collaboration

- Provides opportunities for both institutions to make an impact on pressing global problems through science, research, and education
- Leverages synergies between campus research and MIT LL technology areas to generate innovative solutions
- Exposes a new generation of students to opportunities in engineering, research, and service to the nation and world

Beaver Works is a joint center chartered by the [MIT School of Engineering](#) and [MIT Lincoln Laboratory](#), operated by the Laboratory. Dr. Robert Shin, Head of MIT Lincoln Laboratory's ISR and Tactical Systems Division, is the director of the center. Day-to-day operations are handled by John Vivilecchia, the facility manager. The facility is open to all MIT students, faculty, and collaborators, and provides a nexus for innovation, collaboration, and hands-on development.

The Beaver Works Center was initiated in 2010 through a series of MIT LL-funded capstone research projects in the School of Engineering. In 2013, MIT and MIT LL opened the new, dedicated center designed to facilitate research, workshops, and classwork through the creative fusion of collaborative spaces and prototyping facilities.

ROBUST COMMUNICATIONS FOR AUTONOMOUS SWARMS

Teams of autonomous systems can address key public safety and national security challenges, but building platforms to accomplish these tasks is a highly complex endeavor. Challenges include real-time decision making, coordinated task execution and adaptation, constrained data channels, interference (intentional and unintentional), and scalability. While numerous institutions conduct research into modeling swarm behaviors for application to mechanical systems in the physical world, only a few investigate the communications protocols necessary to enable these behaviors.

This CSAIL capstone course utilizes open architecture principles and adapts COTS hardware and open-source software to enable joint research on communication protocols and autonomy algorithms. The end result will be a hardware and software development platform that can be continuously updated,

upgraded, and refined to accelerate research in this area. MIT LL and MITRE have each contributed \$100K in research funding, along with a staff member to support Prof. Dina Katabi. The staff are providing project ideas, practical implementation help, and serving as mentors for the students taking the course.

APPENDIX C

MIT LINCOLN LABORATORY INVOLVEMENT IN MIT UNDERGRADUATE PROGRAMS

Information about each of the MIT undergraduate programs in which MIT Lincoln Laboratory is currently involved, including the MIT VI-A Masters of Engineering Thesis Program, which is specifically for EECS students, is given below.

MIT Undergraduate Practice Opportunities Program (UPOP): <http://upop.mit.edu/>

Since academic year 2009–2010, the Laboratory has been very actively engaged with the MIT Undergraduate Practice Opportunities Program (UPOP), in both recruiting and sponsorship capacities, as well as volunteer efforts during UPOP on-campus events. The Laboratory has posted UPOP-specific jobs to the UPOP sophomore mailing list every fall since 2009. In addition to hiring UPOP student interns every summer since 2010, the Laboratory has attended the UPOP January Networking Luncheons (held twice over the month of January), years 2010–2016. The Laboratory has also sponsored one of the two annual January Networking Luncheons from 2013–2016. MIT Lincoln Laboratory also sponsored the Spring 2014 UPOP “Round Table” Industry Series.

Along with attending recruiting events with UPOP, one or more Laboratory staff have been present to engage and coach students at multiple touch-points throughout the UPOP year-long programming, including one or more staff at UPOP Mock Interview Nights held three times a year (2009–2016); one or more staff attending as a guest speaker during a spring colloquium dinner, years 2010–2013; and one or more Laboratory staff as a guest speaker during a UPOP Industry Round Tables series (2014–2015).

Presently, two UPOP offers have been extended for 2016 and two others are in process.

Note: Dr. Robert Shin spearheaded the MIT Lincoln Laboratory’s ongoing relationship with this program.

MIT VI-A Masters of Engineering Thesis Program: <http://vi-a.mit.edu/>

Since 1969, the Laboratory has been an industry partner of MIT's Department of Electrical Engineering and Computer Science VI-A Master of Engineering Thesis Program, which matches industry mentors with students in their junior year of study who have demonstrated excellent academic preparation and motivation.

Students in the VI-A program spend two summers as paid interns at the Laboratory, participating in projects related to their fields. Then, the students move on to develop their Masters of Engineering theses under the supervision of both Laboratory engineers and MIT faculty. While the students are working on

their theses, they are supported by a research assistantship funded by the Laboratory and MIT. The Laboratory pays half of the students' annual tuition and MIT pays the balance. The Laboratory also pays a monthly stipend determined by the EECS Dept.

Note: Dr. Eric Dauler is the current MIT LL advisor to this program.

MIT Undergraduate Research Opportunities Program (UROP): <http://web.mit.edu/urop/>

Since 1990, the Laboratory has been one of the centers with which undergraduates may partner under MIT's Undergraduate Research Opportunities Program (UROP). UROP cultivates research partnerships between MIT undergraduates and faculty, offering students the chance to work on cutting-edge research and participate in each phase of standard research activity.

A UROP project may be done at any time during the academic year and/or summer for pay or credit and may take place in any academic department or laboratory. The Laboratory employs several paid UROP interns during the summer and two on average during the academic year.

Interns not hired through UPOP or VI-A are typically hired as UROP.

Note: Professor Jeffrey Shapiro is the current MIT LL advisor to this program.

LIST OF ACRONYMS

ACC	Advanced Concept Committee
AES	Advanced Encryption Standard
AFRL	Air Force Research Laboratory
AS	Autonomous Systems
BAA	Broad Area Announcements
C-WMD	Counter-Weapons of Mass Destruction
COTS	Consumer Off-the-Shelf
CPU	Central Processing Unit
CRIBB	Computational Research in Boston and Beyond
CSAIL	Computer Science and Artificial Language Laboratory
DoD	Department of Defense
EAPS	Department of Earth, Atmospheric and Planetary Sciences
EECS	Electrical Engineering and Computer Science, MIT
FHC	Fully Homomorphic Computing
FPGA	Field-Programmable Gate Array
HLT	Human Language Technology
I&W	Indications and Warnings
ICE	Information Computation Exploitation
IP	Intellectual Property
JIT-ST	Just-In-Time Secure Thread
LLSC	Lincoln Laboratory Supercomputing Center

LNS	Laboratory for Nuclear Science
LOCKMA	Lincoln Open Cryptographic Key Management Architecture
LSP	Lincoln Scholars Program
MGHPCC	Massachusetts Green High-Performance Computing Center
MIT LL	MIT Lincoln Laboratory
ML	Microelectronics Laboratory
NSA	National Security Agency
NTI	New Technology Initiative
ONR	Office of Naval Research
ORAM	Oblivious RAM
PI	Principal Investigator
PUF	Physical Unclonable Function
RHIB	Rigid Hull Inflatable Boat
RLE	Research Laboratory of Electronics
S-COP	Security Coprocessor
SASS	Secure Application-Specific Systems
SLS	Spoken Language Systems
SWaP	Size, Weight, and Power
TCD	Time-Critical Decisions
TO	Technology Office
UAV	Unmanned Aerial Vehicle
UPOP	Undergraduate Practice Opportunities Program
UROP	Undergraduate Research Opportunities Program